

**DYNAMIC COMMAND AND/OR ADDRESS MIRRORING SYSTEM AND  
METHOD FOR MEMORY MODULES**

**TECHNICAL FIELD**

[001] This invention relates to memory modules having memory devices mounted on opposite surfaces of a substrate, and more particularly to memory modules having memory devices on opposite surfaces of a substrate.

**BACKGROUND OF THE INVENTION**

[002] Semiconductor devices, such as memory devices, are normally in the form of a semiconductor substrate or chip mounted in a hermetically sealed package. An integrated circuit fabricated on the chip is then coupled to terminals that are accessible from outside the package. These externally accessible terminals can assume many forms, such as pins projecting outwardly and then downwardly along opposite sides of the integrated circuit package and terminal pads arranged in a group on the bottom of the integrated circuit package, which is known as a ball grid array, or "BGA" configuration.

[003] Each externally accessible terminal of the integrated circuit package is normally associated with a particular function. For example, in an integrated memory device, a first set of externally accessible terminals are input terminals for respective memory address bits A<sub>0</sub>-A<sub>N</sub>, and a second set of externally accessible terminals are input terminals for respective command or status signals C<sub>0</sub>-C<sub>N</sub>, such as RAS\*, CAS\*, and a clock signal. A third set of externally accessible terminals are input/output terminals for respective data bits D<sub>0</sub>-D<sub>N</sub>. Finally, a fourth set of externally accessible terminals are reserved for power and ground.

[004] Although integrated circuits are commonly used singly in many applications, other types of integrated circuits are most commonly used in groups. For

example, memory devices in general, and dynamic random access memory (“DRAM”) devices in particular, are commonly used in groups as part of memory modules. Memory modules are generally in the form of an insulative substrate, such as a printed circuit board, having several memory devices mounted on one or both surfaces of the substrate. Conductors couple the memory devices to connectors that are generally formed by terminals extending along an edge of the substrate. One common memory module is a single in-line memory module, known as a “SIMM,” which includes a single row of memory devices extending across one or both surfaces of the substrate. Another common memory module is a double in-line memory module, known as a “DIMM,” which includes two rows of memory devices extending across one or both surfaces of the substrate.

[005] A common phenomena associated with memory modules, whether SIMMs, DIMMs or some other variety, is the need for an ever increasing storage capacity. For this reason, the capacity of the memory devices mounted on the substrate, as well as the number of externally accessible terminals needed to address the memory devices, has continuously increased. The need for increased memory capacity also increases the need for memory modules having a larger number of memory devices. As a result, memory devices are now usually mounted on both sides of a memory module substrate, and the spacing between memory devices has continued to decrease. The decreased spacing between memory devices and the increased number of terminals has made it more difficult to route conductors to the externally accessible terminals of the memory devices.

[006] One technique that has been used successfully to route conductors to a large number of terminals of closely spaced memory devices is to use substrates having a large number of layers on which conductors are formed. However, it is relatively expensive to provide substrates having a large number of layers, and a large number of closely spaced layers can result in excessive cross-talk between conductors on different layers and excessive conductor capacitance.

[007] Another technique that has made it easier to route conductors to memory device terminals is mirroring in which the terminals of each memory device mounted on one surface of the substrate are positioned directly opposite corresponding terminals of a memory device mounted on the opposite surface of the substrate. This mirroring can occur horizontally, in which corresponding terminals are at the same locations on opposite sides of the mirrored packages, or vertically, in which corresponding terminals are at the same locations above and below a line extending across and bisecting the mirrored packages. In either case, mirroring has the advantage of allowing conductors to extend to a single respective location on the substrate, and to then connect to a respective terminal on each surface of the substrate at that location. Significantly, there is no need to route a conductor coupled to a terminal of an integrated circuit on one surface of the substrate to a different location for coupling to the corresponding terminal of an integrated circuit on the other surface of the substrate.

[008] Although memory device mirroring has the advantage of allowing more compact routing of conductors to the memory devices, it is not without some disadvantages. For memory device mirroring to occur, two different integrated circuit packages must be developed so that corresponding terminals of the two packages are mirror images of each other. The two different packages can theoretically use the same integrated circuit chip, but, in practice, this is not always feasible. In particular, it is important that the integrated circuits on one surface of the substrate respond to signals in the same manner as the integrated circuits on the opposite surface of the substrate. For the circuits to respond in the same manner, it is important for the lengths of corresponding signal paths of the two circuits be identical. Not only is it sometimes difficult to route signal lines from a circuit node to either of two different terminals, doing so creates an undesirable stub connection to the signal path between the terminal and the circuit node. This stub connection can produce signal reflections that can degrade the performance of the integrated circuit. For this reason, it can be necessary to fabricate two different integrated circuit chips, which are the mirror images of each other, for placement in the respective mirrored packages. The need to develop and

stock two different integrated circuit packages, even if the same chip can be used for both packages, can significantly increase the cost of mirrored integrated circuits.

[009] To alleviate the above-described problems of mirroring integrated circuits, programmable integrated circuits have been developed. With reference to Figure 1, an integrated circuit memory device 10 includes a large number of terminals, although only terminals 12, 14 for RAS and CAS signals are shown. The RAS and CAS signals are horizontally mirrored, as explained above. The terminals 12, 14 are each coupled to a respective input of two multiplexers 16, 18. The output of the multiplexer 16 is coupled to a RAS signal node 20, and the output of the multiplexer 18 is coupled to a CAS signal node 22. The multiplexers 16, 18 are controlled by a signal line coupled to an external terminal 26 at a predetermined location. The terminal 12 is coupled to a first input of the multiplexer 16 and to a second input of the multiplexer 18. The terminal 14 is coupled to a second input of the multiplexer 16 and to a first input of the multiplexer 18. As a result, a low applied to the terminal 26 causes the terminal 12 to be coupled to the RAS signal node 20, and the terminal 14 to be coupled to the CAS signal node 22. A high applied to the terminal 26 causes the terminal 12 to be coupled to the CAS signal node 22, and the terminal 14 to be coupled to the RAS signal node 20.

[010] In operation, two of the integrated circuit memory devices 10a,b are mounted on opposite surfaces of a substrate as shown in Figures 2A and 2B, respectively. As a result, the RAS signal is coupled to the terminal 12 of the memory devices 10a and the terminal 14 of the memory device 10b. The CAS signal is coupled to the terminal 14 of the memory devices 10a and the terminal 12 of the memory device 10b. However, the terminal 26 of the memory device 10a is coupled to ground potential, and the terminal 26 of the memory device 10b is coupled to a supply voltage. Therefore, the multiplexer 16 (Figure 1) couples the RAS signal to the RAS signal node 20 of both memory devices 10a,b, and the multiplexer 18 couples the CAS signal to the CAS signal node 22 of both memory devices 10a,b.

[011] The technique explained with reference to Figures 1 and 2A,B has the advantage of allowing mirroring to occur using a single integrated circuit mounted on opposite sides of a substrate, and avoids many of the above-mentioned disadvantages of using two different integrated circuits. However, mirroring using an internal routing circuit, such as the multiplexers 16, 18 shown in Figure 1, has the disadvantage of requiring that a routing circuit for each terminal be fabricated on a semiconductor substrate, thereby using area that could be used for the integrated circuit itself. As a result, the use of routing circuits can significantly increase the cost of memory devices, particularly in view of the large number of terminals present in memory devices that each require a routing circuit, as well as the large number of memory devices included in many systems. The routing circuits can also introduce undesirable delays in the coupling of the RAS and CAS signals to their respective nodes 20, 22.

[012] Another problem in routing conductors to memory devices in memory modules occurs when the memory module includes a memory hub or register through which signals are routed to and from the memory devices. As shown in Figure 3, a memory module 30 includes memory hub 32 mounted on a substrate 34. The memory module 30 also includes a plurality of memory devices mounted on the substrate 34, two of which 38, 40 are shown in Figure 3. In the memory module 30 of Figure 3, each signal transmitted and received by the memory hub 32 is transmitted and received on a first set of terminals coupled to the memory device 38 on the left of the substrate 34 and a second set of terminals coupled to the memory device 40 on the right side of the substrate 34. One of the signals transmitted by the memory hub, *i.e.*, the A<sub>0</sub> address bit, is shown in Figure 3, and this address bit is coupled to correspondingly positioned terminals of the memory devices 38, 40. However, since the A<sub>0</sub> terminal is located on the left side of both memory devices 38, 40, the path to the A<sub>0</sub> terminal of the left memory device 38 is longer than the path to the A<sub>0</sub> terminal of the right memory device 40. As a result, the performance of the two memory devices 38, 40 may not be symmetrical. A similar problem exists when coupling signals from between memory devices and a register (not shown) of a registered memory module.

[013] The above-described difficulties incurred in coupling signals to and from integrated circuits, such as memory devices, creates a need for a mirroring technique that allows a single integrated circuit to be mounted on opposite surfaces of a substrate with correspondingly positioned terminals coupled together, and which does not require internal routing circuitry in each memory device.

#### SUMMARY OF THE INVENTION

[014] A memory module according to one aspect of the invention includes an insulative substrate; on which a plurality of identical memory devices mounted on first and second opposed surfaces of the insulative substrate. The memory devices are mounted on the substrate in a mirrored configuration. As a result, a plurality of terminals of each of the memory devices mounted on the first surface are interconnected to respective, correspondingly positioned terminals of a respective one of the memory devices mounted on the second surface. Address and command signals are coupled to the interconnected terminals. Significantly, either the address signals, the command signals or both the address and command signals are coupled to the interconnected terminals for a plurality of the memory devices in a first configuration if the memory devices mounted on the first surface of the substrate are being accessed. In contrast these signals are coupled to the interconnected terminals for a plurality of the memory devices in a second configuration that is different from the first configuration if the memory devices mounted on the second surface of the substrate are being accessed. The memory devices in one or more memory modules may be coupled to a memory controller that alters the configuration of the signals as indicated above. Alternatively, the configuration of the signals may be altered by either a memory hub or a register mounted on each of the memory modules.

#### BRIEF DESCRIPTION OF THE DRAWINGS

[015] Figure 1 is a block diagram of a conventional integrated memory device that can be used in a memory module in mirrored fashion.

[016] Figures 2A and 2B are terminal diagrams showing the routing of some of the signals for the memory device of Figure 1 when they are mounted on opposite surfaces of a memory module substrate.

[017] Figure 3 is a plan view of a conventional memory module including a memory hub coupled to memory devices on opposite sides of the memory hub.

[018] Figure 4 is a block diagram of a computer system including several memory modules according to one embodiment of the invention.

[019] Figure 5 is a block diagram of one embodiment of a memory controller that may be used in the computer system of Figure 4.

[020] Figure 6 is a block diagram of one embodiment of a computer system according to the present invention.

[021] Figure 7 is a block diagram of still another embodiment of a computer system according to the present invention.

#### DETAILED DESCRIPTION OF THE INVENTION

[022] A computer system 50 according to one embodiment of the invention is shown in Figure 4. The computer system 50 includes a processor 54 for performing various computing functions, such as executing specific software to perform specific calculations or tasks. The processor 54 includes a processor bus 56 that normally includes an address bus, a control bus, and a data bus. The processor bus 56 is typically coupled to a cache memory 58, which, is typically static random access memory (“SRAM”) device. Finally, the processor bus 56 is coupled to a system controller 60, which is also sometimes referred to as a bus bridge.

[023] The system controller 60 contains a memory hub controller 62 that is coupled to the processor 54. The memory hub controller 62 is also coupled to several memory modules 64a-n, and the memory modules 64a-n are coupled to each other, through a downstream bus 66 and an upstream bus 68, which couple data, address and/or control signals away from or toward, respectively, the memory hub controller 62. Each of the memory modules 64a-n includes a memory hub 76 mounted on a substrate

78. The memory hub 76 is coupled to several memory devices 80a,b on one side of the memory module 64 through a first set of command, address and data buses 82, and to several memory devices 86a,b on the other side of the memory module 64 through a second set of command, address and data buses 88. The memory devices 80a,b and 86a,b are identical to each other. The memory hub 76 efficiently routes memory requests and responses between the memory hub controller 62 and the memory devices 80, 86. Computer systems employing this architecture can have a higher bandwidth because the processor 54 can access one memory module 64a-n while another memory module 64a-n is responding to a prior memory access. For example, the processor 54 can output write data to one of the memory modules 64a-n in the system while another memory module 64a-n in the system is preparing to provide read data to the processor 54. The operating efficiency of computer systems using a memory hub architecture can make it more practical to vastly increase data bandwidth of a memory system. A memory hub architecture can also provide greatly increased memory capacity in computer systems.

[024] The system controller 60 also serves as a communications path to the processor 54 for a variety of other components. More specifically, the system controller 60 includes a graphics port that is typically coupled to a graphics controller 90, which is, in turn, coupled to a video terminal 92. The system controller 60 is also coupled to one or more input devices 94, such as a keyboard or a mouse, to allow an operator to interface with the computer system 50. Typically, the computer system 50 also includes one or more output devices 96, such as a printer, coupled to the processor 54 through the system controller 60. One or more data storage devices 98 are also typically coupled to the processor 54 through the system controller 60 to allow the processor 54 to store data or retrieve data from internal or external storage media (not shown). Examples of typical storage devices 98 include hard and floppy disks, tape cassettes, and compact disk read-only memories (CD-ROMs).

[025] With further reference to Figure 4, the memory devices 80a and 86a are mounted on a first surface 100a of the substrate 78, and the memory devices 80b and

86b are mounted on a second surface 100b of the substrate 78. The memory devices 80a, 86a are preferably mounted directly opposite the memory devices 80b, 86b, respectively, and their adjacent terminals are coupled to each other and to signal lines of the buses 82, 88, respectively. As a result, a terminal located at the upper left-hand corner of one of the memory devices 80a is coupled to a correspondingly positioned terminal located at the upper right-hand corner of the opposing memory device 80b, for example.

[026] The terminals of the memory devices 80, 86 are preferably arranged so that data bus terminals of the memory devices 80a, 86a are coupled to data bus terminals of the memory devices 80b, 86b, respectively, on the opposite surface of the substrate 78. Data signals applied to the interconnected terminals of the memory devices will therefore be written to the memory devices 80a, 86a mounted on the first surface 100a as data bits that are different from the data bits written to the memory devices 80b, 86b mounted on the second surface 100b. For example, a data signal for data bit D<sub>1</sub> may be written to a memory device 80a as a D<sub>1</sub> bit and written to a memory device 80b as a D<sub>15</sub> bit. However, when the same data signals are read from the memory devices 80a and 80b, the data signal will be coupled to the signal line for the D<sub>1</sub> data bit even though the data signal was read from the D<sub>15</sub> bit of the memory device 80b. Therefore, as long as the data bus terminals of the memory devices 80a, 86a are coupled to the data bus terminals of the memory devices 80b, 86b, respectively, data will be properly written to and read from the memory devices 80, 86, assuming that all of the data bus terminals to which the data bits are coupled are associated with the same data strobe signal.

[027] Unlike the data bus terminals of the memory devices 80, 86, the address and control bus terminals of the memory devices 80, 86 are not interchangeable. Therefore, an address signal coupled to a terminal of the memory device 80a cannot be simply coupled to a correspondingly positioned terminal of an opposing memory device 80b. Instead, each address and control signal must be coupled to a specific terminal of each of the memory devices 80, 86 regardless of the location of the memory devices.

Rather than using the conventional approaches of using different memory devices having different terminal configurations or routing signals to the proper signal nodes, the address and control signals are coupled to the proper terminals of the memory devices 80, 86 by the memory hub 76 coupling different signals to the same lines of the buses 82, 88 depending upon which memory device 80, 86 is being accessed. For example, if the A<sub>0</sub> address bit of the memory device 80a is connected to the A<sub>10</sub> address bit of the opposing memory device 80b, the memory hub 76 will couple an address signal for bit A<sub>0</sub> to a specific signal line of the bus 82 if the memory device 80a is being addressed, and it may couple an address signal for bit A<sub>10</sub> to the same signal line of the bus 82 if the memory device 80b is being addressed. Conversely, the memory device will couple an address signal for bit A<sub>0</sub> to one signal line of the bus 82 if the memory device 80a is being addressed, and it will couple the address signal for bit A<sub>0</sub> to another signal line of the bus 82 if the memory device 80b is being addressed. Different address and control signals are also coupled to the same lines of the bus 88 depending on whether memory device 86a or 86b is being accessed.

[028] The memory devices 80, 86 may be either of two types of memory devices each of which operates in either of two modes. Furthermore, the memory devices 80, 86 have different terminal assignments depending on which mode of operation is used by the memory devices 80, 86. For example, as is known to one skilled in the art, DRAM memory devices operating in one double data rate (“DDR”) mode known as the “DDR2” mode use one set of terminal assignments, and, DRAM memory devices operating in another DDR mode known as the “DDR3” mode, use another set of terminal assignments. For example, the same terminal to which an A<sub>5</sub> address bit is coupled to memory devices operating in the DDR2 mode will receive an A<sub>12</sub> address bit for memory devices operating in the DDR3 mode. In each case, the memory hub 76 will couple the proper signal to each signal line of each of the buses 82, 88 depending on: (1) whether the memory devices 80, 86 are of the type that operate in the DDR2 mode or the DDR3 mode; (2) whether the memory devices 80, 86 are to the left of the memory hub 76 and are thus coupled to the hub 76 through the bus 82 or to

the right of the memory hub 76 and are thus coupled to the hub 76 through the bus 88; and (3) whether the memory devices 80, 86 are mounted on the first surface 100a of the substrate 78 or the second surface 100b of the substrate 78.

[029] In one embodiment of the invention, the memory hub 76 will couple signal to the memory devices 80, 86 according to Table 1 below in which A<sub>0</sub>-A<sub>16</sub> are row and column addresses, BA<sub>0</sub>-BA<sub>3</sub> are bank addresses, CASZ is an active low column address strobe signal, CSZ0 and CSZ1 are active low chip select signals to select memory devices 80, 86 on the first surface 100a and second surface 100b, respectively, ODT0 and ODT1 are signals, RASZ is an active low row address strobe signal and WEZ is an active low write enable signal.

Table 1

Hub Left Terminal Location	Hub Right Terminal Location	82a, 86a DDR2 CS <sub>0</sub>	82b, 86b DDR2 CS <sub>1</sub>	82a, 86a DDR3 CS <sub>0</sub>	82b, 86b DDR3 CS <sub>1</sub>
C <sub>5</sub>	C <sub>27</sub>	A <sub>0</sub>	A <sub>10</sub>	A <sub>3</sub>	A <sub>0</sub>
C <sub>11</sub>	B <sub>21</sub>	A <sub>1</sub>	A <sub>2</sub>	A <sub>6</sub>	A <sub>5</sub>
C <sub>6</sub>	C <sub>26</sub>	A <sub>2</sub>	A <sub>1</sub>	BA <sub>3</sub>	BA <sub>2</sub>
C <sub>12</sub>	B <sub>20</sub>	A <sub>3</sub>	A <sub>4</sub>	A <sub>2</sub>	A <sub>1</sub>
B <sub>5</sub>	B <sub>27</sub>	A <sub>4</sub>	A <sub>3</sub>	A <sub>1</sub>	A <sub>2</sub>
B <sub>11</sub>	A <sub>21</sub>	A <sub>5</sub>	A <sub>6</sub>	A <sub>12</sub>	A <sub>13</sub>
B <sub>6</sub>	B <sub>26</sub>	A <sub>6</sub>	A <sub>5</sub>	A <sub>5</sub>	A <sub>6</sub>
B <sub>12</sub>	C <sub>19</sub>	A <sub>7</sub>	A <sub>8</sub>	A <sub>7</sub>	A <sub>8</sub>
A <sub>5</sub>	B <sub>29</sub>	A <sub>8</sub>	A <sub>7</sub>	A <sub>8</sub>	A <sub>7</sub>
A <sub>13</sub>	A <sub>19</sub>	A <sub>9</sub>	A <sub>11</sub>	A <sub>9</sub>	A <sub>11</sub>
D <sub>12</sub>	C <sub>21</sub>	A <sub>10</sub>	A <sub>0</sub>	BA <sub>2</sub>	BA <sub>3</sub>
A <sub>4</sub>	A <sub>29</sub>	A <sub>11</sub>	A <sub>9</sub>	A <sub>11</sub>	A <sub>9</sub>
A <sub>12</sub>	A <sub>20</sub>	A <sub>12</sub>	A <sub>13</sub>	A <sub>4</sub>	A <sub>14</sub>
A <sub>6</sub>	A <sub>27</sub>	A <sub>13</sub>	A <sub>12</sub>	A <sub>13</sub>	A <sub>12</sub>
A <sub>11</sub>	A <sub>22</sub>	A <sub>14</sub>	A <sub>15</sub>	A <sub>15</sub>	A <sub>15</sub>

A <sub>7</sub>	A <sub>28</sub>	A <sub>15</sub>	A <sub>14</sub>	A <sub>14</sub>	A <sub>4</sub>
A <sub>10</sub>	A <sub>23</sub>	A <sub>16</sub>	A <sub>16</sub>	A <sub>16</sub>	A <sub>16</sub>
D <sub>7</sub>	D <sub>27</sub>	BA <sub>0</sub>	BA <sub>0</sub>	BA <sub>1</sub>	BA <sub>0</sub>
E <sub>13</sub>	D <sub>20</sub>	BA <sub>1</sub>	CASZ	A <sub>10</sub>	WEZ
E <sub>7</sub>	E <sub>25</sub>	BA <sub>2</sub>	BA <sub>2</sub>	CSZ <sub>1</sub>	CSZ <sub>1</sub>
F <sub>12</sub>	F <sub>22</sub>	BA <sub>3</sub>	BA <sub>3</sub>	ODT <sub>1</sub>	ODT <sub>1</sub>
D <sub>8</sub>	D <sub>26</sub>	CASZ	BA <sub>1</sub>	CASZ	RASZ
D <sub>6</sub>	D <sub>28</sub>	CSZ <sub>0</sub>	CSZ <sub>0</sub>	WEZ	A <sub>10</sub>
D <sub>11</sub>	D <sub>21</sub>	CSZ <sub>1</sub>	CSZ <sub>1</sub>	BA <sub>0</sub>	BA <sub>1</sub>
C <sub>9</sub>	E <sub>28</sub>	ODT <sub>0</sub>	ODT <sub>0</sub>	ODT <sub>0</sub>	ODT <sub>0</sub>
D <sub>13</sub>	C <sub>20</sub>	ODT <sub>1</sub>	ODT <sub>1</sub>	A <sub>0</sub>	A <sub>3</sub>
F <sub>9</sub>	E <sub>22</sub>	RASZ	WEZ	CSZ <sub>0</sub>	CSZ <sub>0</sub>
E <sub>10</sub>	D <sub>22</sub>	WEZ	RASZ	RASZ	CASZ

[030] In one embodiment, the memory hub 76 has externally accessible terminals arranged in a grid, and the location of each terminal is designated by a letter designating the row of the terminal and a number designating the column of the terminal. For example, A1 is a terminal located at the upper left hand corner of the memory hub 76, and F29 is a terminal located at the lower right hand corner of the memory hub 76. The locations of these terminals that are coupled to the memory devices 80 through the bus 82 are listed in the first column of Table 1, and the locations of these terminals that are coupled to the memory devices 86 through the bus 88 are listed in the second column of Table 1. The signals applied to the terminals of the memory hub 76 when the memory devices 82a, 86a on the first surface 100a of the substrate 78 are being accessed and the devices 82a, 86a are of the type operating in the DDR2 mode are listed in the third column of Table 1. The signals applied to the terminals of the memory hub 76 when the memory devices 82b, 86b on the second surface 100b are being accessed and the devices 82b, 86b are of the type operating in the DDR2 mode are listed in the fourth column of Table 1. The signals applied to the

terminals of the memory hub 76 when the memory devices 82a, 86a are of the type operating in the DDR3 mode and are being accessed are listed in the fifth column of Table 1. Finally, the signals applied to the terminals of the memory hub 76 when the memory devices 82b, 86b are of the type operating in the DDR3 mode and are being accessed are listed in the sixth column of Table 1.

[031] For example, the C5 terminal of the memory hub 76, which is coupled to a signal line of the bus 82, receives the A<sub>0</sub> address bit when the memory devices 80a are being accessed and they are of the type operating in the DDR2 mode, and it receives the A<sub>3</sub> address bit when the memory devices 80a are being accessed and they are of the type operating in the DDR3 mode. If the memory devices 86a are being accessed, the C27 terminal of the memory hub 76, which is coupled to a signal line of the bus 88, receives the A<sub>0</sub> address bit when the memory devices 86a are being accessed and they are of the type operating in the DDR2 mode, and it receives the A<sub>3</sub> address bit when the memory devices 80a are being accessed and they are of the type operating in the DDR3 mode. If the memory devices 80b are being accessed, the C5 terminal of the memory hub 76, which is coupled to a signal line of the bus 82, receives the A<sub>10</sub> address bit when the memory devices 80b are of the type operating in the DDR2 mode, and it receives the A<sub>0</sub> address bit when the memory devices 80b are being accessed and they are of the type operating in the DDR3 mode. If the memory devices 86b are being accessed, the C27 terminal of the memory hub 76, which is coupled to a signal line of the bus 88, receives the A<sub>10</sub> address bit when the memory devices 80b are of the type operating in the DDR2 mode, and it receives the A<sub>0</sub> address bit when the memory devices 80b are being accessed and they are of the type operating in the DDR3 mode.

[032] The routing of signals to the terminals of the memory devices 80, 86 based on whether they are of the type operating in the DDR2 or the DDR3 mode is static since it does not change during the operation of the memory devices 80, 86. However, the routing of signals from either the left set of terminals A1-F13 of the memory hub 76 to the memory devices 80 or from right set of terminals A19-29 of the

memory hub 76 and to the memory devices 86 is dynamic, and therefore changes rapidly while the memory devices 80, 86 are being accessed.

[033] Although the address and most of the control signals must be dynamically mirrored, in one embodiment some control signals can be applied to multiple symmetrically positioned terminals of the memory device 80, 86 so that dynamic mirroring is not required. For example, a first clock enable CKE signal is applied to a terminal of the memory device 80, and a second clock enable CKE signal is applied to a terminal of the memory device 86 regardless of whether the memory devices 80, 86 are of the type operating in the DDR2 or the DDR3 mode. Regardless of which memory device 80, 86 is being addressed, the addressed memory device will receive a CKE signal at the appropriate terminal. The signals for which dynamic mirroring is not required, as well as the locations of the terminals of the memory hub 76 to which those signals are applied in one embodiment, are listed in Table 2, below.

Table 2

Hub Left Terminal Location	Hub Right Terminal Location	82a, 86a DDR2 CS <sub>0</sub>	82b, 86b DDR2 CS <sub>1</sub>	82a, 86a DDR3 CS <sub>0</sub>	82b, 86b DDR3 CS <sub>1</sub>
D <sub>9</sub>	E <sub>19</sub>	CKE <sub>0</sub>	CKE <sub>0</sub>	CKE <sub>0</sub>	CKE <sub>0</sub>
F <sub>6</sub>	E <sub>29</sub>	CKE <sub>1</sub>	CKE <sub>1</sub>	CKE <sub>1</sub>	CKE <sub>1</sub>
F <sub>8</sub>	F <sub>26</sub>	CLK <sub>0</sub>	CLK <sub>0</sub>	CLK <sub>0</sub>	CLK <sub>0</sub>
G <sub>11</sub>	G <sub>24</sub>	CLK <sub>1</sub>	CLK <sub>1</sub>	CLK <sub>1</sub>	CLK <sub>1</sub>
F <sub>7</sub>	F <sub>25</sub>	CLKZ <sub>0</sub>	CLKZ <sub>0</sub>	CLKZ <sub>0</sub>	CLKZ <sub>0</sub>
G <sub>12</sub>	G <sub>23</sub>	CLKZ <sub>1</sub>	CLKZ <sub>1</sub>	CLKZ <sub>1</sub>	CLKZ <sub>1</sub>
F <sub>13</sub>				RESET	RESET

[034] The memory hub 76 used in the memory modules 64a-n of Figure 4 includes a memory controller, and one embodiment of a memory controller 200 that can be used is shown in Figure 5. Briefly, the memory controller 200 receives high-level macro commands, such as ACTIVATE ROW, COLUMN, and PRECHARGE, converts these commands to DRAM commands, schedules the DRAM commands for outputting

at the proper time, and routes the scheduled DRAM commands to the correct terminals of the memory hub 76, as previously explained. With reference to Figure 5, the memory controller 200 includes a command queue 204 that receives the high-level macro commands from the memory hub controller 62 (Figure 4). The Command Queue 204 translates the received macro commands to DRAM command signals, such as RASZ, CASZ, WEZ, etc., places the command signals in a queue in the order that their corresponding macro commands were received, and then outputs the commands in the proper order. The Command Queue 204 also receives address signals, which are placed in a queue in the order that they were received, and they are subsequently output in the proper order. The command and address signals that are output from the Command Queue 204 are the command and address signals listed in Tables 1 and 2.

[035] The DRAM command signals and address signals from the Command Queue 204 are applied to a Command Scheduler 210, which spaces the command and address signals apart from each other with the proper delay. The delay is measured in periods of a clock CLK signal, which is also applied to the Command Scheduler 210. For example, the Command Scheduler 210 might schedules the CASZ to be output three clock periods after the RASZ signals was output from the Command Scheduler 210. When the command and address signals are output from the Command Scheduler 210, they are stored in a micro Command Shifter 214, which is basically a shift register that is driven by the CLK signal. The command and address signals are stored in the Command Shifter 214 in the proper order and with the proper spacing because they were shifted into the Command Shifter 214 from the Command Scheduler 210 in that manner.

[036] When the properly timed and ordered command and address signals are shifted out of the Command Shifter 214 responsive to the CLK signal, they are applied to one input of a Multiplexer 220. The Multiplexer 220 also receives the DRAM command and address signals at a second input directly from the Command Queue 204. In operation, the Multiplixer 220 couples the command and address signals to its output directly from the Command Shifter 214 as these signals are received when the memory

devices 80, 86 have been idle. As a result, the latency penalty that would be incurred in being coupled through the Command Scheduler 210 and the Micro Command Shifter 214 is avoided. After the initial command and address signals have been coupled to the output of the Multiplexer 220, the Multiplexer 220 selects the output of the Micro Command Shifter 214 for coupling to its output.

[037] The command and address signals at the output of the Multiplexer 220 are applied to Signal Swap Multiplexers 230 that are controlled by a CONFIG command coupled to the Multiplexers 230 through a bus 232. The CONFIG command indicates whether the memory devices 80, 86 are of the type operating in the DDR2 or DDR3 modes. The Multiplexers 230 are also controlled by CSZ<sub>0</sub> and CSZ<sub>1</sub> signals, which indicate whether the memory devices 80a, 86a on the first surface 100a or the memory devices 80b, 86b on the second surface 100b are being accessed. The Signal Swap Multiplexers 230 can be implemented using a large number of individual multiplexers arranged in a matrix in a manner that will be apparent to one skilled in the art. The Signal Swap Multiplexers 230 route the command and address signals to the terminals of the memory hub 76 as shown in Tables 1 and 2.

[038] The properly routed and timed command and address signals are coupled from the Signal Swap Multiplexers 230 to Ring Buffers 240, 242. The Ring Buffer 240 is coupled through the bus 82 (Figure 4) to the memory devices 80 on the left hand side of the memory modules 64, and the Ring Buffer 242 is coupled through the bus 88 to the memory devices 86 on the right hand side of the memory modules 64. Each of the Ring Buffers 240, 242 is basically a first-in, first-out buffer that is driven by one of the clock signals CLK<sub>0</sub>, CLK<sub>1</sub>, CLKZ<sub>0</sub> or CLKZ<sub>1</sub> that is applied to the memory devices 80, 86. The command and address signals are therefore shifted from the Ring Buffers 240, 242 and applied to the terminals of the memory hub 76 in synchronism with the operation of the memory devices 80, 86.

[039] Although the memory controller 200 can be used in the memory hub 76 in the computer architecture shown in Figure 4, it can also be used as a stand-alone memory controller in the computer system 300 shown in Figure 6. The computer

system 300 uses many of the same components that are used in the computer system 50 of Figure 4. Therefore, in the interests of brevity, these components have been provided with the same reference numerals, and an explanation of their operation will not be repeated. The computer system 300 differs from the computer system 50 by using memory modules 310 that do not contain a memory hub. Instead, the memory devices 80, 86 in each memory module 310 are coupled directly to the memory controller 200 in the system controller 60 through a memory bus 320. The memory controller 200 applies the command and address signals shown in Tables 1 and 2 to the signal lines of the bus 320 depending on which memory devices 80a, 86a, 80b or 86b are being accessed and whether the memory devices 80, 86 are of the type operating in either the DDR2 mode or the DDR3 mode, as explained above and as shown in Table 1.

[040] Another embodiment of a computer system 400 according to the invention is shown in Figure 7. The computer system 400 also uses many of the same components that are used in the computer systems 50 and 300 of Figure 4 and 6, respectively, and an explanation of the operation of these components will not be repeated. The computer system 400 differs from the computer system 300 by using memory modules 410 having registers 420 coupling the command and address signals to the memory devices 80, 86. Registered memory modules are well known in the art. The computer system 400 uses a memory controller 430 in the system controller 60 that may operate in the same manner as the memory controller 200 to apply the command and address signals to the registers 420 depending on which memory devices 80a, 86a, 80b or 86b are being accessed and whether the memory devices 80, 86 are of the type operating in either the DDR2 mode or the DDR3 mode. Alternatively, the memory controller 420 may operate in a conventional manner, and circuitry like the Signal Swap Multiplexers 230 shown in Figure 5 may be included in each register 420 to apply the command and address signals to the memory devices 80, 86 depending on which memory devices 80a, 86a, 80b or 86b are being accessed and whether the memory devices 80, 86 are of the type operating in either the DDR2 mode or the DDR3 mode. Other variations will be apparent to one skilled in the art.

[041] The various embodiments of the inventions have the advantage of routing the proper command and address signals to the memory devices 80, 86 at the memory controller level rather than at each of the memory device level. As a result, the number of circuits needed to perform this function is markedly less than if such circuitry was in each of the memory devices 80, 86. The various embodiments thus allow a mirroring with single memory device used on both surfaces of the substrate and on both the left and right sides of the substrate. From the foregoing it will be appreciated that, although specific embodiments of the invention have been described herein for purposes of illustration, it will be understood by one skilled in the art that various modifications may be made without deviating from the spirit and scope of the invention. For example, instead of using registered memory modules, circuitry like the Signal Swap Multiplexers 230 could simply be used on each memory module. Accordingly, the invention is not limited except as by the appended claims.